

[illegible]

5 This application is related to, and claims the benefit of, U.S. Provisional Patent Application Serial No. 60/176,411, entitled "System and Method for Effective and Efficient Electronic Communication Management," filed January 13, 2000, which is hereby incorporated by reference. This application is also related to U.S. Patent Application No. 09/602,588, entitled "System and
10 Method for Automatic Task Prioritization," filed June 21, 2000, and U.S. Patent Application No. 09/624,361, entitled "System and Method for Optimizing Timing of Responses to Customer Communications," filed July 24, 2000, which are hereby incorporated by reference.

15

1.

20

2.

1

channels such as letters and telephone calls, customers may also communicate with an organization via electronic mail, facsimile, web-based forms, web-based chat, and wireless communication and voice. An organization will most likely incorporate these and any other newly developed communication channels to

5 allow customers to communicate in a way they find most convenient.

Many of the communication channels mentioned above contain information that is unstructured in nature, usually expressed in natural language. Different customers may make identical requests each in a unique way, using different communication channels, different words, or both.

10 Human agents are usually required to review each natural language communication to evaluate the customer's intent, and to determine what information or action would be responsive to that intent.

Agents typically must look to various sources to gather all of the information required to respond appropriately to a customer communication.

15 The information may be retrieved from a variety of sources, such as legacy systems, databases, back office systems, and front office systems. Each of these sources may store data in a unique structure or format. An agent typically gathers and organizes the required information from one or more of these information sources and uses the information to compose an appropriate

20 content-rich reply that is responsive to the customer's intent.

Utilizing people to respond to customer communications is often rather inefficient. In addition, an increase in the number of communications received

by an organization typically requires an exponential increase in the number of people required to provide an acceptable level of customer service.

Several types of automatic systems exist for responding to customer communications. Rule-based systems, keyword-based systems, and statistical systems typically do not perform with the necessary accuracy to substantially automate business processes, such as responding to customer inquiries, and require a large investment in resources to keep them up-to-date. Many learning systems utilize a training set of data that is a poor representation of the system's world, which reduces the accuracy of the system and makes the process of updating the system very cumbersome.

SUMMARY OF INVENTION

The invention provides a system and method for electronic communication management. The system comprises a contact center, a modeling engine, an adaptive knowledge base, and a feedback module. The contact center may send and receive communications via various communication channels including phone, facsimile, electronic mail, web forms, chat, and wireless. The modeling engine analyzes received communications to determine an intent. For received communications containing natural language text, the modeling engine performs morphological, semantic, and other analyses. For voice-based communications, the system performs various digital signal processing tasks.

The adaptive knowledge base stores models that are used to predict responses and actions to the received communications based on the intent identified by the modeling engine. The feedback module monitors actual responses to the received communications and compares them to the predicted responses. If a predicted response is substantially the same as the actual response, the model or models that predicted the response are updated with positive feedback. The feedback module supports multiple feedbacks to a single communication. If a predicted response is substantially different than the actual response, the model or models that predicted the response are updated with negative feedback. The feedback process may be performed either in real time or off-line. Each model has an internal accuracy gauge that

is updated by the feedback. The system learns from every communication that is processed.

The modeling engine may also support various application specific modules, for example, an automatic response module, an automatic task prioritization module, an expertise based routing module, a content filter, a workflow application module, and a business process automation module. The modeling engine may also retrieve data from various sources, such as databases and back office systems, which relate to the intent of a communication.

The contact center converts each received communication into a universal data model format. The models in the adaptive knowledge base may also be expressed in the universal data model format, so that models of different types of data may be compared to each other.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram for one embodiment of an electronic communications system, according to the present invention;

5

FIG. 2 is a block diagram for one embodiment of the Modeling Engine of FIG.1, according to the present invention;

FIG. 3 is a flowchart of method steps for processing relationship events,
10 according to one embodiment of the present invention; and

FIG. 4 is a diagram of relationship event processing, according to one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Referring now to FIG. 1, a block diagram of an electronic communication management system 100 is shown. System 100 includes, but is not limited to, a contact center 112, a universal data model 114, a modeling engine 116, an adaptive knowledge base 118, and data access services 120. Contact center 112 receives communications from a variety of channels. The channels include, but are not limited to, telephone 130, facsimile 131, electronic mail (Email) 132, web-based communications 133, chat communications 134, and wireless communications 135. Other types of electronic communications 136, for example a file transfer using the File Transfer Protocol (FTP), are within the scope of the present invention. New communication channels may be added without taking the system off-line.

The communications received by contact center 112 typically contain data or information that is unstructured in nature. With the exception of some web-based or other type of electronic forms, the communications usually contain information expressed in natural language. Each individual correspondent may compose a communication in a unique way, even when requesting the same type of information.

System 100 also includes data access services (middleware) 120 that retrieve data from various sources that include, but are not limited to, legacy systems 122, front office systems 123, back office systems 124, and databases 125, 126. These sources typically contain data that are structured, each

source potentially having its own data format. Data access services 120 of the present invention translate the data from each source to conform to a universal data model (UDM) 114 format, so that data from the various sources may be represented according to a common data structure. For example, a universal
5 data model of front office data will have the same structure as a model of back office data.

Contact center 112 of the present invention translates incoming communications from the various communication channels to conform to universal data model 114, so that data from the various communication
10 channels may be represented according to a common data structure. For example, a universal data model of a facsimile communication will have the same structure as a model of an email communication with any obviously necessary changes. The universal data model 114 of the present invention allows different types of data to be represented in a common data format,
15 regardless of the source or type of data.

UDM 114 allows system 100 to analyze, model, and compare models of different types of data. System 100 may create and compare models of email communications, models of database information, and models of human
agents. System 100 is able to compare a model of an email communication to
20 a model of an agent, which in turn may be compared to a model of a business process because all models have the same universal structure. The universal data model 114 advantageously allows previously incompatible types of

information to be processed by the same system. The UDM 114 is a data driven model of information.

In one embodiment of system 100, UDM 114 includes concepts that are produced in a hierarchical processing scheme. Lower level concepts may be as general as single words from email texts, voice data, or may be as specific as field descriptors from a web-based form. Further processing allows system 100 to infer intents and other higher level concepts from the lower level concepts. Thus, a final representation of information in UDM 114 is identical for all information sources.

System 100 also includes a modeling engine (ME) 116. ME 116 is a tool that, in conjunction with UDM 114, allows system 100 to perform a myriad of tasks using data from various sources. ME 116 supports various Application Specific Modules (ASM) 140-146. The ASMs perform specialized operations using ME 116 in conjunction with models and information in the universal data format, and are discussed in further detail below.

To be able to support the ASMs, ME 116 monitors relationship events and business processes, and looks for semantical and other patterns. Relationship events are any communications between the organization and other external or internal entities. Relationship events may, for example, include an email from a customer, an order placed via a secure web-based ordering system, an email sent from a queue to an agent, a document

submitted to a document management system, and an automatic email response sent to a customer.

ME 116 builds semantical models based on relationship events. ME 116 continuously updates the models using positive and negative feedback (explicit
5 or implicit) from system 100, and stores the models in Adaptive Knowledge Base (AKB) 118. Using the models, ME 116 learns the business processes of the system. The models in AKB 118 have strong predictive powers. Given a relationship event, a model is able to predict which other relationship events are likely to follow. For example, given an inquiry, system 100 can predict
10 what is the most likely business response, such as answer, routing, or data association.

ME 116 monitors what business processes follow from a relationship event and determines whether these processes match a model's predictions. A response that matches the prediction is positive feedback that increases the
15 model's accuracy rating, and a response that does not match the prediction is negative feedback that decreases the model's accuracy rating. Both positive and negative feedback adapts AKB 118. Feedback in the present invention is further discussed below in conjunction with FIG. 4.

ME 116 knows when it knows, and knows when it doesn't know, based
20 on measured correlations between confirmed and unconfirmed predictions. ME 116 analyzes a relationship event and is able to recognize when the relationship event corresponds to a model in AKB 118 and when the event does not correspond to a model, or corresponds to a low-dependability model. When

a relationship event does not correspond to a model, ME 116 will typically route the event for handling by an agent, and then create a potential new model based on the event, or use predefined rules.

Since received communications are translated into the universal data
5 format, ME 116 is able to learn from feedback for one communication channel and apply that knowledge to another communication channel. For example, feedback from an agent responding to inquiries received via Email 132 may result in knowledge that allows system 100 to automatically answer inquiries on a chat 134 channel, without reconfiguration of system 100.

10 AKB 118 organizes the models into various categories. Logically related categories are associated with a branch, which in turn may be associated with a branch of larger scope. Using similar statistical techniques to the ones described, creation of the hierarchies can be either manual (via a configuration tool or API) or automatic by monitoring feedback. In addition, some branches
15 may be created with associated rules, which allows system 100 to be fine tuned and to detect non-business-compliant agent actions (e.g., submission of a secure communication over an unsecured communication channel).

AKB 118 may also include flat hierarchies as a special case of tree
hierarchies. Other types of graphs, such as a cyclic layered graph, are within
20 the scope of the invention. Incorporating relationship events into the branches of AKB 118 improves the predictive confidence of the branches. As new relationship events are received, new categories are created and new branches develop in AKB 118. The models are used to identify the underlying intent of a

relationship event, and to categorize the event based on various criteria, some of which are manual (based on rules) and others which are automatic (based on learning or training). AKB 118 categorizes events using a meta-language that is able to combine manual and automatic criteria.

5 ME 116 creates branches in AKB 118 using a semi-automatic process. At system start-up, the branching process is given information regarding the tasks or applications system 100 is to support, and some general guidelines on how to differentiate between them. The guidelines typically include hard rules combined with intent-based rules.

10 When system 100 is up and running, ME 116 uses feedback to modify the branches in AKB 118. ME 116 collects statistical data for intent-based branches, alerts system 100 when hard rules are violated, and monitors the performance of AKB 118. ME 116 also suggests structural changes to AKB 118. ME 116 may join branches that have similar statistical profiles. ME 116
15 may split branches into sub-branches using a modified expectation maximization process to increase the overall performance of system 100. ME 116 may also perform all of the above operations on a flat structure without branches.

20 Referring now to FIG. 2, a block diagram of ME 116 is shown, according to one embodiment of the present invention. ME 116 includes, but is not limited to, a Natural Language Processor (NLP) 210, and a Statistical Modeler 212. Natural Language Processor 210 evaluates customer communications in

natural language to determine the customer's intent and other relevant information. Data structures such as fixed fields in a web-based form are considered a subset of natural language.

5 NLP 210 includes, but is not limited to, a Language Dependent Module (LDM) 220 that extracts information from a natural language communication, and a Lexical Knowledge Base (LKB) 222 that include lexical, morphological, and semantic information. NLP 210 may identify the language of the communication and have separate LDMs 220 and dictionaries for various languages that operate simultaneously. LDM 220 identifies the content-rich
10 parts of the relationship event, and corrects spelling and grammatical errors. In fact, LDM 220 expects poor grammar and broken text in relationship events, and does not rely upon accurate grammar to perform sentence analysis, which improves the accuracy of system 100.

LDM 220 performs text analysis using context rules. Some context rules
15 are rigid rules and others are learned statistically from actual texts collected from previous relationship events. Grammar errors and broken text have only a "local" influence on the entire text analysis. Common errors are learned and are referred to as non-rigid rules.

LDM 220 may be modified to parse and understand specific classes of
20 grammatical and syntactic errors that are characteristic of speakers of a particular non-native language. For example, a native German speaker writing in English may connect adjectives together as one word. LDM 220 will

recognize this type of error as one commonly made by native German speakers, and correct the error accordingly.

NLP 210 also collects other types of information about a relationship event. This information may include the use of passive voice, semantic
5 information, contextual information, morphological information, and
quantitative information. The quantitative information may include the
number of sentences or number of exclamation points. NLP 210 identifies key
concepts in a relationship event, which are basic components of language
information. These components include, but are not limited to, sentences,
10 words, stems, semantically significant constructs, the type and number of
punctuation marks, strong use of the passive voice, dates, and currency
amounts.

NLP 210 also identifies the modality of text from any text channel. NLP
210 identifies modes that include, but are not limited to, emotive language, the
15 age or education of the writer, indications of large sums of money, and
offensive content. NLP 210 also may identify the type of communication, for
example inquiry, request, complaint, formal letter, resume, joke
communication, and junk mail.

NLP 210 also includes Lexical Knowledge Base (LKB) 222. LKB 222
20 includes lexical, morphological, and semantic domains. The information in
LKB 222 may be either general or domain dependent. LKB 222 may include,
but is not limited to, a lexicon, a thesaurus, a spell checker, a morphological
analyzer, and a tagger.

LKB 222 may be constructed off-line using predetermined semantical data, and may accumulate semantical data on-line by monitoring user actions. The semantic domain of LKB 222 is the main resource for ME 116 in creating meaningful models.

5 NLP 210 builds semantic models for relationship events using concepts described in or otherwise suggested by the event and the relationships between the concepts. As a very simple example, the statement "The book is on the table" contains the concepts "book" and "table." A semantic model would include these two concepts and show their relationship of one being "on" the
10 other. NLP 210 also generalizes concepts based on feedback. For example, the statement "the book is on the table" could be generalized to a model that describes a written object on a piece of furniture. Thus, the statement "the newspaper is on the chair" could result in the same generalized model.

The second main component of ME 116 is the Statistical Modeler 212.
15 Modeler 212 is used in training the system and creating models in conjunction with NLP 210. Statistical modeler 212 performs relationship algebra using the models in Adaptive Knowledge Base 118. Relationship algebra compares and modifies models.

For example, model A and model B represent business processes. If
20 model A is compared to model B, a statistical score may be 70%. "Adding" model A with model B produces a new model A' ($A+B=A'$). If model A' is compared to model B, the statistical score may be 72%. By combining model B with model A, the resulting model A' is a little more like model B. "Subtracting"

model B from model A produces another new model A'' (A-B=A''). If model A'' is compared to model B, the statistical score may be 68%. Thus model A'' is a little less like model B. ME 116 uses these and other types of relationship algebra operations to continuously update models. Feedback determines
5 which operations are performed with the models. System 100 may expect some erroneous feedback, so not every operation necessarily directly affects the models. In some cases, ME 116 will look for supportive evidence before a particular action is incorporated into a model.

ME 116 maintains internal queues of potential models and potential
10 concepts that are not in actual usage by system 100. ME 116 continuously checks and updates these potential models and concepts. ME 116 automatically collects domain knowledge, which is information about the world in which it lives. ME 116 creates the potential, or hypothetical, models when it sees a pattern of relationship events in the system. ME 116 tests these
15 hypothetical models and incorporates data into the models when the data is deemed sufficiently accurate.

There are two potential main sources for loss of accuracy of models in a modeling system. The first source is variance, where there is not enough data to support a model. The second source is bias, where there are false beliefs
20 about the data. ME 116 is capable of distinguishing between these two sources for loss of accuracy, and is able to accurately assess the amount of data that supports each model. ME 116 may thus decide when the data is sufficiently rich to support predictive use of a model.

Each model in Adaptive Knowledge Base 118 has an internal accuracy gauge that is updated continuously by feedback from the system. Positive feedback increases a model's accuracy rating, and negative feedback decreases the model's accuracy rating. Each model compares its prediction with the actual result or action of the system and responsively modifies its accuracy rating accordingly.

A model's accuracy is rated by recall and precision. A model's recall is the ratio of the number of events the model identified to the number of events the model should have identified. A model's precision is the ratio of the number of events the model identified correctly to the number of events the model identified. Recall and precision may be traded against one another. For example, high recall can be achieved by indiscriminately identifying all events, however this results in loss of precision. Alternatively, high precision can be achieved by identifying events for which the model has high confidence, but some events may be missed and recall would be lower. A good model should have high recall and high precision. Recall and precision measurements may be assessed using feedback from system 100. A statistical matching value between documents and models may also be evaluated by a calculated statistical likelihood value. The likelihood value may be calculated using an adaptive kernel method based on match value results of various categories.

For each branch, AKB 118 gathers statistical data that distinguishes the branch model from its background using significant concepts. AKB 118 contains two types of data, active concepts and candidate concepts that may

become active in the future. Feedback from system 100 is used to update both types of data. Each concept has an associated rate that relates to the difference between its expected contribution in relevant and irrelevant events.

The feedback process modifies the rates of the concepts. Each newly
5 modified rate determines whether a candidate concept should become active, and whether an active concept should become inactive. The contribution of an active concept to branch classification is based on a log-likelihood-ratio between two statistical models that are built by interpolating the statistical data of the relevant models and the irrelevant models.

10 If a model's accuracy rating drops below a predetermined threshold, the model is marked as inaccurate and is not used by the system to make decisions. However, the model still receives data and makes predictions, which are compared to the actual performance of the system. The model continues to receive feedback and the accuracy rating continues to be updated. Use of the
15 model will be resumed if and when the model regains an accuracy rating above the predetermined threshold.

When a model is being disregarded, the system may behave according to some predetermined rules, use keywords, or perform some other action like sending relationship events to a queue for processing by a human agent.

20 As ME 116 learns the business processes of the system, it becomes able to identify erroneous or malicious input by an agent. An agent may make errors that unintentionally damage the accuracy of the system, or may intentionally take incorrect actions to deliberately sabotage the accuracy of the

system. ME 116 can identify responses made by agents that do not fit the patterns predicted by the models. ME 116 identifies these responses by analyzing feedback that is very different from the system's predictions.

Suspicious responses are identified according to the reliability of the origin of
5 the response, and the difference between the system's decision and the feedback. ME 116 analyzes feedback according to its origin and will trace suspicious feedback, then obtain verification before using the feedback.

There are several ways in which ME 116 may be trained when a system initially goes on-line at an organization. In one embodiment, ME 116 is placed
10 on-line with no active models. The system then receives live relationship events and begins building models. The accuracy of the models' predictions will increase as the number of relationship events increases.

In another embodiment, ME 116 receives historical relationship event data and builds models based on this data. Thus, ME 116 may be placed on-
15 line with models in place that will be fairly accurate if the historical relationship event data is a fair representation of live relationship events.

In a further embodiment, ME 116 is placed on-line with no active models, and the system behaves according to a set of rules or logical expressions. ME 116 builds models based on live relationship events while the
20 rules are in place. When the models reach an acceptable level of accuracy, the rules are automatically disregarded and the models take over the processing of events. If a model becomes inaccurate, the system may default back to the rules until the model becomes accurate again.

Returning to FIG. 1, one of the Application Specific Modules supported by ME 116 is an Automatic Response (AR) module 140. This module automatically composes and sends an appropriate pre-written or "canned" response to a customer communication. Exemplary responses may contain general information about the organization, a request for more information from the customer, or a confirmation that the communication was received. A related response module may compose relevant content-rich responses to customer communications using fragments of existing text and retrieved data, with or without the involvement of human agents. For example, an automatic response concerning a customer's account balance may contain a pre-existing text message with the appropriate amount inserted into the message. An automatic response may also be a combination of multiple fragments identified by ME 116. ME 116 analyzes the customer's message to identify intents and/or categories. ME 116 then fetches data, for example an account balance or order status, and a pre-existing text message associated with the appropriate category.

Another module is an Expertise Based Routing (EBR) module 142 that routes a customer communication to the agent or queue best qualified to perform the required task or solve the customer's problem. EBR 142 compares a model of the relationship event (e.g., customer facsimile or wireless communication) with models of all available agents or queues to determine which agent is best suited for responding to the event, and then routes the

event to that agent or queue. An agent's model may include, but is not limited to, the agent's seniority, automatically detected areas of competency, and languages.

Automatic Task Prioritization (ATP) module 141 is another module that is
5 supported by ME 116. ATP 141 prioritizes tasks and learns from feedback and rules which tasks have a higher priority than others. Priority may be determined based on a service level agreement with the particular customer, dollar amounts mentioned in the communication, the type of information being requested by the customer, or other content of a customer communication.
10 ATP 141 may be customized to make priority decisions according to an organization's specifications.

ME 116 also may support a content filter module 143 that filters responses composed by agents. Content filter 143 may be used to avoid emotive or rude replies to customers, and as a method for assessing the quality
15 of the responses. Content filter 143 may also be used to ensure compliance with an organization's regulations. In another embodiment, content filter 143 may filter customer communications for emotive or offensive content, and route these communications to a queue for handling by a senior agent.

Business Process Automation (BPA) modules 145 may be used to
20 complete routine business processes automatically. For example, a transfer of funds between a customer's accounts in a financial institution may be handled automatically, or monthly shipments of a product from a retailer may be processed automatically for customers with standing orders. An automatic

business process may be performed using data retrieved from various sources internal to the organization. Thus, a legacy system and a database having different data structures may exchange data via a BPA module 145 supported by ME 116.

5 Other modules that may be supported by ME 116 are Workflow Applications (WFA) 144, which allow agents to escalate relationship events, reassign events, or add comments to messages. ME 116 may support other ASMs 146 that an organization may require, such as automatic fetching of data and/or agents needed to solve a particular problem, building a team of skilled
10 agents according to the needs of a task, suggesting internal recipients for a communication within an organization, and detecting "hot prospects."

Another ASM may automatically generate relevant Frequently Asked Questions (FAQ) that are responsive to a customer's intent. The module, in conjunction with ME 116, determines the intent or intents of the customer,
15 selects from a general list the most relevant FAQs, and incorporates them into a response sent to the customer. Yet another module may post an automatic response to a relationship event on a website and send a customized URL to the customer so that, among other things, the system can track if and when the response was accessed.

20 Another Application Specific Module may generally classify communications and documents based on content. Customer communications that may not require a response, for example responses to customer surveys, and other electronic documents, such as directives from management to

agents, may be classified by content and automatically stored in an appropriate file or database. For instance, this module may identify extremely negative survey responses and forward them to a queue for further evaluation by management.

5 Other Application Specific Modules that an organization desires may be built and incorporated into an existing system without taking the system off-line. Thus, each system may be customized to meet the needs of a specific organization and may be updated and modified as the organization's needs change.

10

Referring now to FIG. 3, a flowchart of method steps for processing a relationship event is shown, according to one embodiment of the present invention. In the FIG. 3 embodiment, the new relationship event is received via a text-based channel such as email.

15

In step 310, system 100 receives a new relationship event, which is translated into a universal data model by contact center 112, and is then routed to ME 116. In step 312, the Natural Language Processor 210 analyzes the event to identify concepts, utilizing linguistic data from Adaptive Knowledge Base 118. NLP 210 may perform various analyses on the event, including
20 semantic, contextual, morphological, and quantitative.

Next, in step 314, the concepts are used to build a model for the event using statistical modeling and modeler 212, as discussed above. In step 316, modeler 212 determines whether it needs further linguistic information for the

event. If so, the method returns to step 312 for additional natural language processing. If not, the method continues with step 318, where ME 116 maps the event model to all models in AKB 118 to determine the relevancy of the event to each category. The event mapping step assigns a score to every
5 category for each relationship event based on how closely the model for the relationship event corresponds to a category's models. The score is determined by comparing the models using the relationship algebra described above. In other embodiments, logical expressions (rules) are used to categorize events. These rules may also be used when models are considered inaccurate.

10 Next, in step 320, the event is routed for automatic or semi-automatic action, based on the category scores and configuration settings. An event may be routed to certain queues or agents if the corresponding category score is greater than a predetermined threshold. The user (manager) of system 100 may set these thresholds and vary them at any time to best suit the needs of
15 the organization. Alternatively, the threshold values may be set automatically based on information from the system itself.

Relationship events received via a voice channel are processed slightly differently. Voice events may be initially handled by an agent who determines the customer's intent. The agent is presented with a tree showing topics of
20 various customer intents. The agent chooses an appropriate topic or topics, and the system then fetches data and canned responses corresponding to the selected topic.

Voice events may also be processed by a digital signal processing (DSP) module that categorizes events based on the acoustical content of an audio signal. The module compares a received voice event to models of previous events to predict an appropriate action, including transmitting a pre-recorded
5 vocal response. Voice events may be processed in real time, or may be stored as voice mail messages and processed off-line. In the preferred embodiment, the voice events are not transformed into text before being categorized. Agent feedback may be used to refine the models of acoustical patterns.

Sub 10
B1
Referring now to FIG. 4, a diagram of relationship event processing is shown, according to one embodiment of the present invention. A relationship event is received in the contact center 410 and translated into the universal data format. The event is then processed by the Modeling engine (ME) 412 in conjunction with the Adaptive Knowledge Base (AKB), as described above in
15 conjunction with FIGS. 2 and 3. ME 412 accesses any required data from data access services 414 and forwards the event model and data for further processing.

Sub 20
B2
ME 412 may forward the event model and data to an automatic response module 416, an assisted response module 418, or a queue 420. The present invention may also include other modules, as described above in conjunction with FIG. 1. ME 412 may forward the event model to as many modules as needed to respond to all of the intents expressed in the event.

Sub
B3

The automatic response module 416 generates an appropriate automatic response and forwards the response to an audit module 424. The audit module 424 may or may not perform an audit on the response, as will be described below. If an audit is performed, the result is then forwarded to a feedback module 426, where feedback is sent to ME 412. This feedback from an automatic response will most likely be positive feedback that strengthens the accuracy rating of the model that selected the response. The automatic response is then sent to the contact center 410, where the response is formatted for the appropriate communication channel and sent to the customer. Feedback module 426 supports multiple feedbacks to a single communication.

Sub
B4

The assisted response module 418 will forward the event model, the associated information gathered by ME 412 including a history of interactions with the customer, and a list of suggested (canned) responses to the event to an agent 422. The agent 422 may select one or more of the suggested responses, or may compose an original response. The response is forwarded to the audit module 424, which may or may not perform an audit on the response. The response then flows to the feedback module 426, which provides the response feedback to ME 412.

The feedback system of the present invention performs two separate processes: updates the structure of models in AKB 118 and adjusts the models' accuracy ratings. The feedback from feedback module 426 may be positive or negative. If the agent selected one of the suggested responses, the model that

predicted that response will be updated and its accuracy rating will increase since its feedback was positive. The models that predicted other responses will also be updated, and their accuracy ratings will decrease since their predictions were not implemented, thus their feedback was negative. If the agent composed an original response, some or all of the models will receive negative feedback.

Sub B5 Relationship events and associated data may be sent to one or more queues 420 by ME 412. Queues may store events of low priority until events of high priority have been processed. Other queues may store events that contain more than one request. For instance, a customer may request information regarding an existing account and express an interest in a new account. The first request may be satisfied with an automatic response, but the second request may be sent to a queue for new accounts. The second request may then be forwarded to an agent who handles new accounts.

Sub B6 The present invention includes built-in quality control based on audits of responses to relationship events. The audit module 424 reviews responses to relationship events and feeds this information back to ME 412 via the feedback module 426. ME 412 may determine that a particular agent assisted response was inappropriate if the response varies greatly from what was predicted. The system user may configure the audit module 424 to perform audits based on various criteria, including, but not limited to, the experience level of agents, the status of the customer based on a service level agreement, which queue the

Be
could

event was routed to, the channel of the event, the type of response, and whether the agent sent a canned or a composed response.

The learning capabilities of ME 116 allow the system to maintain a high level of recall without sacrificing precision. Recall is a ratio of the number of relationship events that are received by the system and the number of events selected for automatic response or automatic action. Precision is a ratio of the number of events selected for automatic response and the number of responses that were correct. In typical systems, when a system is instructed to increase the number of events that will be responded to automatically, the precision of the system decreases noticeably. When recall is increased, the system will select events for automatic response in which the system has lower confidence. This results in a higher potential for errors in selecting appropriate responses, which lowers the system's precision.

In the system of the present invention, ME 116 allows system 100 to automatically respond to a large number of relationship events correctly. ME 116 quickly learns from feedback which responses are appropriate for various intents, and automatically creates new models as new types of relationship events are received. Thus, system 100 may be instructed to increase the number of events selected for automatic response without causing a significant loss of precision.

Loss of precision usually occurs because the "world" a system lives in is continuously changing. A static rule-based or keyword-based system becomes

less accurate over time. In contrast, ME 116 learns and adapts with every relationship event that it sees, thus maintaining a high level of accuracy over time.

The invention has been explained above with reference to a preferred embodiment. Other embodiments will be apparent to those skilled in the art in light of this disclosure. For example, the present invention may readily be implemented using configurations other than those described in the preferred embodiment above. Additionally, the present invention may effectively be used in conjunction with systems other than the one described above as the preferred embodiment. The present invention, which is limited only by the appended claims, is intended to cover these and other variations upon the preferred embodiment.